NWG/RFC# 623
Comments on On-Line Host Name Service

MCK 26-FEB-74 16:21 22004

22004 Postel

RECEIVED
MAY 28 1974

RFC #623
NIC #22004
references: RFC # 606, 608

Mark Krilanovich
UCSB
February 22, 1974

Comments on On-line Host Name Service

Peter Deutch in RFC # 606 pointed out the desireability of having a single host maintain a data base containing official host names and host addresses, as well as other information of secondary importance. Mike Kudlick in RFC # 608 agreed with the concept, and proposed that the NIC would implement Peter's ideas. I would like to add my voice to those in support of such a service, and express a few ideas for its modification.

The notion of having a single host maintain this data base clearly has the weakness that anyone wishing to obtain a copy of the data may be faced with the situation that the serving host is not available when the data is desired. It is true that each host could save a copy of the most recently obtained data, such that whenever a current copy cannot be obtained, at least a very recent copy is available. This is not a particularly attractive idea, since it requires a non-trivial amount of bother on the part of everyone. Therefore, I propose that the NIC maintain the master data base, and one other host be responsible for maintaining a secondary copy, which is to be updated to be equal to the NIC's at periodic and often intervals, such as once a day. This way, anyone wishing to obtain the data can first try the NIC, and if that fails, try the secondary host, thus much reducing the probability that the data cannot be obtained, while requiring additional software to be written at only one additional host. Further, I volunteer UCSB to be that secondary host.

The proposal currently underway calls for the host names data base to have the format of an ASCII file. RFC 606 makes the point, with which I comletely agree, that this data base should be formatted in an easily machine-readable form. To this end, I propose that the data base be retrievable in binary form rather than ASCII. Using this concept, for example, <host-address> would be a one-byte (eight-bit) binary number, <host-name> would be a one-byte length field followed by that many ASCII characters, and the possible <attribute-values>'s for the STATUS <attribute-name> would be one-byte binary numbers. This modification would clearly make the data base unintelligible to a human user, and, just as clearly, much more easily interpretted by a program.

RFC 608 states that the data base will be maintained as a

file, and retrievable through FTP. I question the wisdom of basing such a simple process as keeping a host table up to date on such a complex protocol as FTP. Therefore I propose that the data base be available via a program running under its own socket at the NIC and at the secondary host. This also avoids the necessity for the accessing program to know the login parameters for the guest account at the serving host, which in fact might not be the same at the two hosts. Again, the motivation is to make things easy for accessing programs.

Anyone with comments about any of the above is encouaged to make them known.